

CORRECTED VERSION

(19) World Intellectual Property  
Organization  
International Bureau



(43) International Publication Date  
11 March 2004 (11.03.2004)

PCT

(10) International Publication Number  
WO 2004/021196 A1

(51) International Patent Classification<sup>7</sup>: G06F 12/08

(21) International Application Number:  
PCT/US2003/026128

(22) International Filing Date: 21 August 2003 (21.08.2003)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
60/406,623 28 August 2002 (28.08.2002) US

(71) Applicant (for all designated States except US): GRASS  
VALLEY (U.S.) INC. [US/US]; 400 Providence Mine  
Road, Nevada City, CA 95959 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): ROSKER, Steven,  
Brian [US/US]; 32 Saddle Club Road, Lexington, MA  
02420 (US). SINGER, Charles, Todd [US/US]; 34 Hick-  
ory Lane, North Reading, MA 01864 (US).

(74) Agents: TRIPOLI, Joseph, S. et al.; Thomson Licensing  
Inc., Two Independence Way, Suite #200, Princeton, NJ  
08540 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU,  
AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU,  
CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH,  
GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC,  
LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW,  
MX, MZ, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC,  
SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA,  
UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

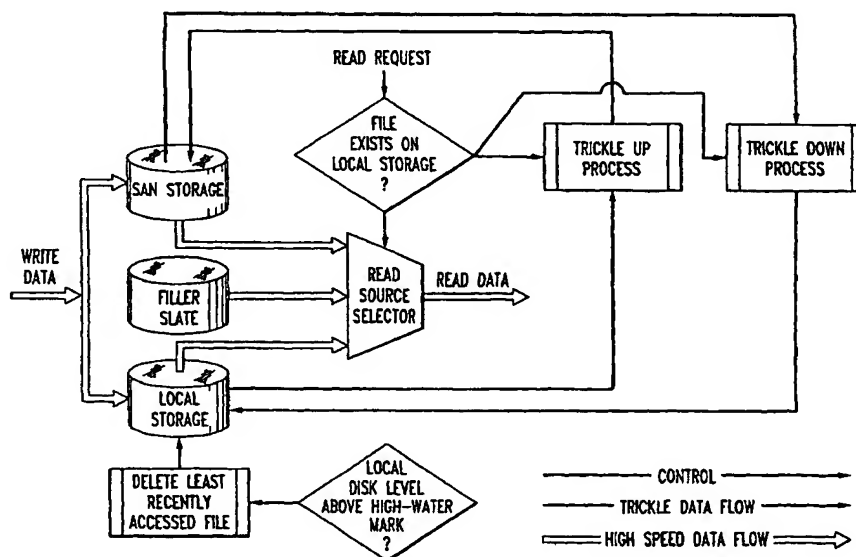
(84) Designated States (*regional*): ARIPO patent (GH, GM,  
KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW),  
Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),  
European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE,  
ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO,  
SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM,  
GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— with international search report

[Continued on next page]

(54) Title: VIDEO-STORAGE NETWORK HAVING INCREASED PERFORMANCE



(57) Abstract: A video storage system (10) includes a storage area network (12) and at least one local cache storage unit (14). Both the storage mechanism and the local cache storage unit store at least some content in common. In response to a read request, a read director (20) determine which of the storage mechanism and local cache storage unit contains the requested content. Upon determining that the requested content resides on the local cache storage unit, the read director directs the content request to the local cache storage unit, thereby reducing the demand on the storage mechanism. If the content does not reside on the local cache storage unit, read director directs the request to the storage mechanism, but if the content is unavailable, the content request will be filled with filler data from a filler data source (24).

WO 2004/021196 A1

WO 2004/021196 A1



(48) Date of publication of this corrected version:

22 July 2004

(15) Information about Correction:

see PCT Gazette No. 30/2004 of 22 July 2004, Section II

*For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

## VIDEO-STORAGE NETWORK HAVING INCREASED PERFORMANCE

### CROSS REFERENCE TO RELATED APPLICATION

5

This application claims priority under 35 U.S.C. 119(e) to U.S. Provisional Patent Application Serial No. 60/406,623, filed August 28, 2002, and incorporated by reference herein.

### TECHNICAL FIELD

10

This invention relates to a video storage network for storing video information.

### BACKGROUND ART

15

In the past, video editing occurred on stand-alone devices, typically videotape machines. An editor seeking to create a video segment comprised of separate clips would queue the clips in the desired order on one or more source machines and then select desired clip for recording. The advent of shared video storage greatly changed the manner in which video editing now occurs. With present day shared storage networks, multiple users can each edit a copy of the same

20 content using non-linear editing techniques. Many present-day video storage networks often place limits on the number of users because of bandwidth constraints. The switch fabric of such video storage networks often will not allow more than a limited number of users to access the storage facility. Overcoming this drawback by providing additional links increases overall cost.

25

Thus, there is need for a technique for providing a shared video storage system that overcomes the aforementioned disadvantage.

### BRIEF SUMMARY

30

Briefly, in accordance with a preferred embodiment, there is provided a storage system for storing content, and preferably, video files. The storage system of the present principles includes a storage mechanism, a large disk drive, or one or more Redundant Arrays of Inexpensive Disk Drives (RAIDs). The storage system further includes at least one local cache storage unit for mirroring at least a portion of the content stored on the storage mechanism. Thus, a user can obtain such mirrored content from the local cache storage unit rather than the

storage mechanism, thus reducing the bandwidth demands on the storage mechanism. A write director is coupled to the storage mechanism and the one local cache storage unit for controlling the incoming content written into the storage mechanism and the local cache storage unit. A cache manager manages the content coherency between the storage mechanism and the local  
5 cache storage unit so that at least some of the content on the storage mechanism also exists on the local cache storage unit and vice versa. A read director responds to a content request from a user and directs the request to one of the storage mechanism and the local cache storage unit depending on the content availability of each. In this way, if the content is available on the local cache storage unit, the user reads from that unit rather than the storage mechanism. A storage  
10 mechanism access manager monitors the reading of content from, and the writing of content to the storage mechanism to control the read and write managers accordingly.

#### BRIEF SUMMARY OF THE DRAWING

15 FIGURE 1 depicts a block schematic diagram of a storage system in accordance with a preferred embodiment of the present principles.

#### DETAILED DESCRIPTION

20 FIGURE depicts a storage system 10 in accordance with a preferred embodiment of the present principles that affords improved efficiency. At its heart, the storage system 10 includes a storage mechanism 12 comprised of a single disk drive, or one or more Redundant Arrays of Inexpensive Disk Drives (RAIDs) that receive content for subsequent access by a user. Typically, the storage mechanism 12 stores video files with or without embedded audio for  
25 editing and/or broadcast. Having all users access the storage mechanism 12 at any given time will constrain bandwidth, thus imposing limitations on the maximum number of users. To overcome this disadvantage, the storage system 10 of the present principles also includes at least one local cache storage unit 14, and preferably, a plurality of such units. Each local cache storage unit 14 typically takes the form of a disk drive or one or more RAID's.

30 As discussed in greater detail below, the local cache storage unit 14 mirrors at least some of the content stored on the storage mechanism 12. In other words, the local cache storage unit 14 stores at least some of the same files as the storage mechanism 12. In this way, a user seeking to access a particular file stored both on the storage mechanism 12 and on the local cache storage

unit 14 can access the local cache storage unit rather than the storage mechanism. Accessing the local cache storage unit 14 reduces bandwidth demand on the storage mechanism 12. In practice, different groups of users each can have a dedicated local cache storage unit. Alternatively, assignment of a local cache storage unit 14 to a group of users can occur on a dynamic basis.

5       A write director 16 controls the writing of input content from an external source (not shown) to each of storage mechanism 12 and the local cache. Such input content can include live video received from a camera or a satellite feed, or previously recorded video. In practice, the write director 16 takes the form of a programmed processor, hard-wired logic circuit, or field programmable array (FPGA) or the like that responds to control information received from a  
10       storage mechanism access manager 17 that monitors the read and write load on the storage mechanism 12. In practice, the storage mechanism Access Manager 17 takes the form of a programmed processor, hard-wired logic circuit, or field programmable array (FPGA). In accordance with the output signal from the storage mechanism Access Manager 17, the write director 16 directs incoming content to either or both of the storage mechanism 12 and the local  
15       cache storage unit 14.

      A cache manager 18, typically in the form of a programmed processor, hard-wired logic circuit, or field programmable array (FPGA) or the like, controls the writing of content between the local cache storage unit 14 and the storage mechanism 12 in accordance with control signals from the local cache storage unit to maintain content coherency. In other words, the cache  
20       manager 18 maintains at least some content in common on the storage mechanism 12 and the local cache storage unit 14. To accomplish this task, the cache manager 18 undertakes background copying of files to the local cache storage unit 14 from the storage mechanism 12 ("trickle down"). Likewise, the cache manager 18 undertakes background copying of files to the storage mechanism 12 from the local cache storage unit 14 ("trickle up"). Such trickle up  
25       becomes necessary when a user edits content for subsequent use by others. The rate at which "trickling down" and "trickling up" occurs is a function of the overall bandwidth demand on the storage mechanism 12. The cache manager 18 also manages the storage capacity of the local cache storage unit 14 by deleting infrequently accessed files. Should the available storage capacity of the local cache storage unit 14 exceed a prescribed level, the cache manager 18 will  
30       delete the least recently accessed file first. Thereafter, the cache manager 18 will delete the next least recently accessed file and so on until the available storage capacity of the local cache storage unit 14 reaches an acceptable level.

To read content from the storage system 10, a user enters a read request received by a read director 20. In practice, the read director 20 takes the form of a programmed processor that monitors the content on both the storage mechanism 12 and the local cache storage unit 14. Upon receipt of a read request, the read director 20 determines whether the requested content  
5 resides on the local cache storage unit 14. If so, the read director 20 sends a command to a read source selector 22, typically a multiplexer, to designate the local cache storage unit 14 as the source of the requested content. Thus, under such circumstances, the read director 20 "redirects" the read request to the local cache storage unit 14. In the event the local cache storage unit 14 lacks the requested content, the read director 20 causes the source selector 22 to select the storage  
10 mechanism 12. Alternatively, the read director 20 will cause read source selector 22 to select a storage unit 24 provide the user with filler content in the event the local cache storage unit 14 can not provide the requested content and insufficient bandwidth exists to permit access to the storage mechanism 12. When the local cache storage unit 14 lacks the requested content, the read director 20 alerts the cache manager 18. In response, the cache manager 18 initiates a trickle  
15 down copy operation to transfer the content from the storage mechanism 12 to the local cache storage unit 14.

In operation, the storage system 10 advantageously offers improved performance by reducing bandwidth demands on the storage mechanism 12. The local cache storage unit 14 serves to mirror at least some of the content on the storage mechanism 12, thus providing an  
20 alternate source for supplying such content to a user. The cache manager 18 carefully manages the local cache storage unit 14 in response to the state of the storage mechanism 12, as monitored by the storage mechanism access manager 17, to assure that local cache storage unit will contain the requisite content. The storage mechanism access manager controls the reading of content from, and the writing of content to the storage mechanism during intervals of high bandwidth  
25 demand.

The foregoing describes a storage network having improved performance.

CLAIMS

- 1           1.     A storage system, comprising:  
2           a storage mechanism for storing content;  
3           at least one local cache storage unit for mirroring at least a portion of the content stored  
4           on the storage mechanism;  
5           a write director coupled to the storage mechanism and to the at least one local storage  
6           cache for controlling content written into the storage mechanism and to the at least one local  
7           storage cache;  
8           a cache manager for managing content copying between the storage mechanism and the at  
9           least one local storage cache to maintain at least partial content coherency;  
10          a read director responsive to a request for content from a user for directing said content  
11          request to a selected one of the at least one local storage cache and the storage mechanism  
12          depending on content availability of each; and  
13          a storage mechanism access manager for monitoring read and write loading of the storage  
14          mechanism and for controlling the read and write directors and the cache manager in accordance  
15          with the storage mechanism read and write loading.
- 1           2.     The storage system of claim 1 wherein the cache manager manages the storage  
2           capacity of the local cache storage unit by successively deleting a least accessed file until the  
3           local cache storage unit has an available storage capacity above a prescribed level.
- 1           3.     The storage system of claim 1 further comprises a filler storage unit for storing  
2           filler content, and wherein the read director directs the read request to the filler storage unit to  
3           provide filler content when the requested content is unavailable from the storage mechanism and  
4           the local storage cache unit.
- 1           4.     The storage system of claim 1 wherein the storage mechanism further comprises  
2           at least one disk drive.
- 1           5.     The storage system of claim 1 wherein the storage mechanism further comprises  
2           at least one Redundant Array of Inexpensive Disk Drives (RAID).

1           6.     The storage system of claim 1 wherein the local cache storage unit further  
2 comprises at least one disk drive.

1           7.     The storage system of claim 1 wherein the local cache storage unit further  
2 comprises at least one Redundant Array of Inexpensive Disk Drives (RAID).

1           8.     The storage system according to claim 1 wherein the read director redirects a  
2 request for content to the local cache storage unit when the requested content is available thereat  
3 to reduce bandwidth requirements on the storage mechanism.

1           9.     The storage system according to claim 1 wherein the cache manager copies at  
2 least some content from the storage mechanism to the local cache storage unit previously  
3 unavailable on the local cache storage unit.

1           10.    The system according to claim 1 wherein the storage mechanism access manager  
2 controls the read and write directors to reduce reading from, and writing to the storage  
3 mechanism during intervals of limited storage mechanism bandwidth.

1           11.    A method for storing content, comprising the steps of:  
2           writing incoming content to at least one of a Storage Area Network (storage mechanism)  
3 and a local cache storage unit;  
4           monitoring content coherency between the storage mechanism and the local cache storage  
5 unit;  
6           copying content between the storage mechanism and the local cache storage unit in  
7 accordance with the content coherency therebetween  
8           directing a request for content from a user to a selected one of the storage mechanism and  
9 the local cache storage unit depending on the content availability of each,  
10          monitoring read and write loading of the storage mechanism; and  
11          controlling reading of content from, and writing of content to the storage mechanism in  
12 accordance with the storage mechanism read and write loading.



1           12.     The method according to claim 11 wherein the step of directing the content  
2 request further comprises re-directing the content request to the local cache storage unit if the  
3 requested content resides at the local cache storage unit.

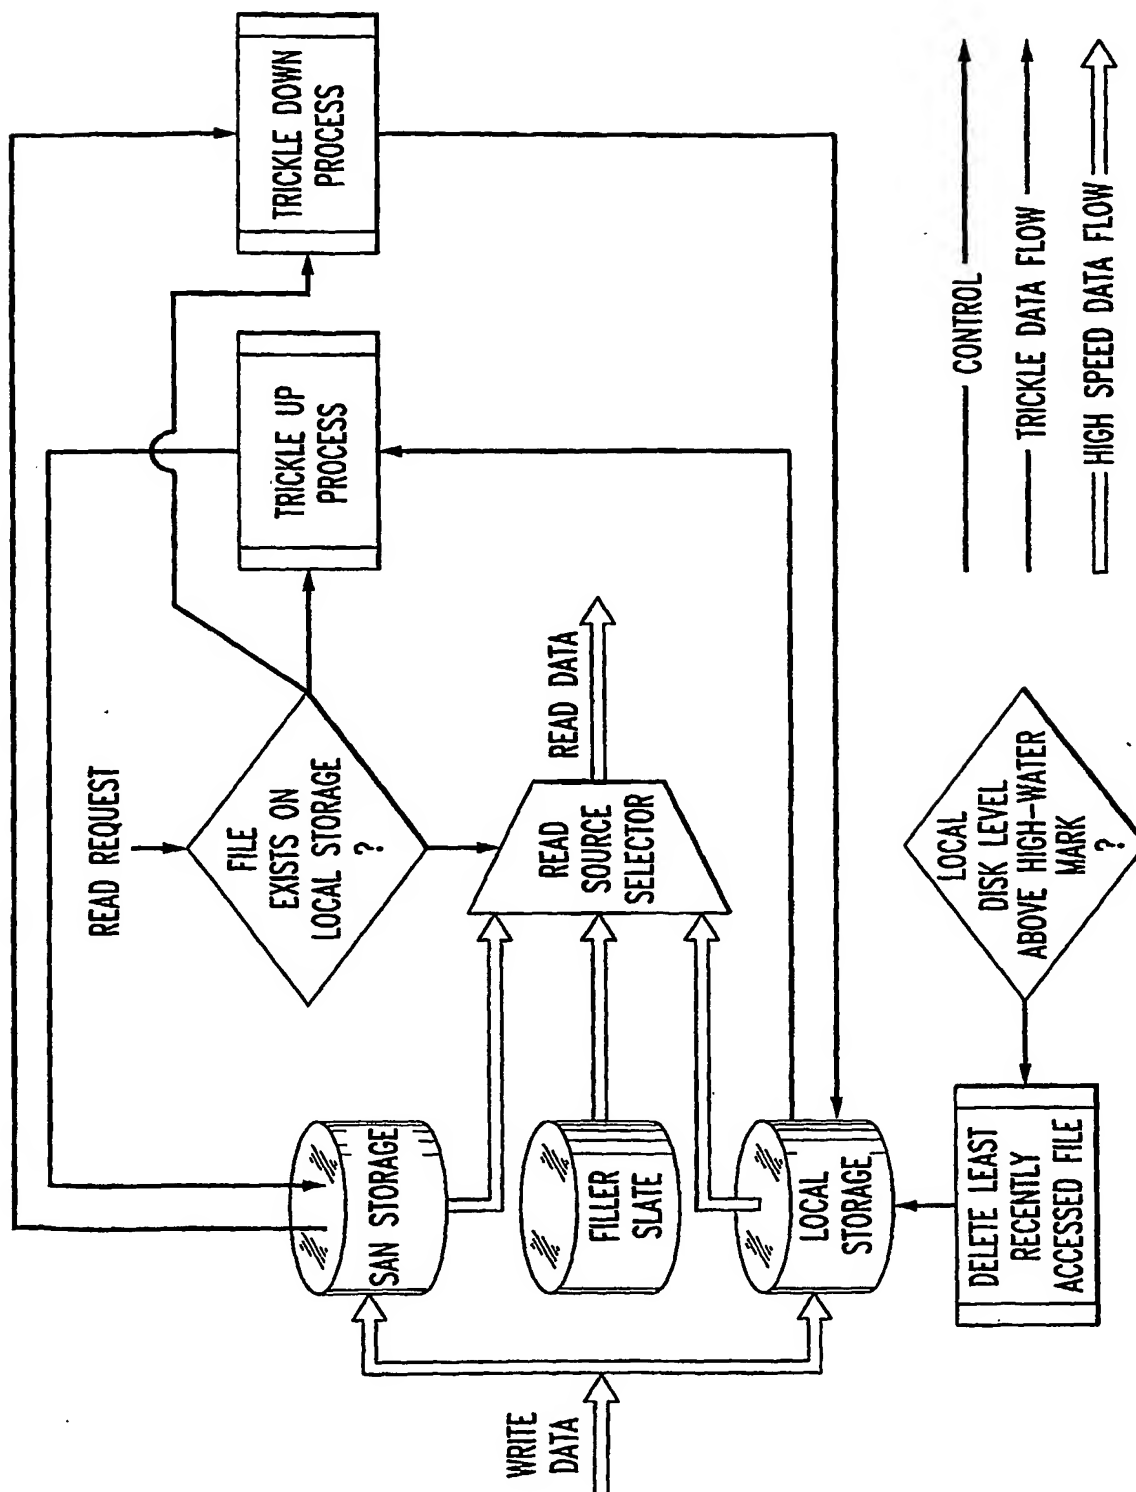
1           13.     The method according to claim 11 wherein the step of directing the content  
2 request further comprises re-directing the content request to a filler storage unit to provide filler  
3 if the requested content neither resides at the local cache storage unit or at the storage  
4 mechanism.

1           14.     The method according to claim 13 further comprising the step of writing content  
2 from the storage mechanism to.

1           15.     The method according to claim 11 further comprising the step of writing content  
2 from the local cache storage unit to the storage mechanism.

1           16.     The method according to claim 11 wherein the step of controlling reading of  
2 content from, and writing of content to the storage mechanism further comprises the step of  
3 restricting access to the storage mechanism during intervals of high bandwidth demand.

1/1



# INTERNATIONAL SEARCH REPORT

International application No.

PCT/US03/26128

## A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : G06F 12/08  
US CL : 711/118, 141

According to International Patent Classification (IPC) or to both national classification and IPC

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)  
U.S. : 711/118, 141

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)  
IEEE xplore

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X, P — Y, P	US 2002/0194429 A1 (Chiu et al.) 19 December 2002 (19.12.2002), Figure 3, Abstract, Page 3, Paragraphs 0037-0043, Page 1, Paragraph 0012	1-2, 4-12, 15, 16 3, 13
A, P	US 6,463,509 B1 (Teoman et al.) 08 October 2002 (8.10.2002)	1-16
Y, P	US 2002/0184460 A1 (Tremblay et al.) 05 December 2002 (5.12.2002), Psge 2, Paragraph 0027	3, 13

☐ Further documents are listed in the continuation of Box C.

☐ See patent family annex.

\* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T"

later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X"

document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y"

document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&"

document member of the same patent family

Date of the actual completion of the international search

25 November 2003 (25.11.2003)

Date of mailing of the international search report

15 DEC 2003

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US  
Commissioner for Patents  
P.O. Box 1450  
Alexandria, Virginia 22313-1450

Facsimile No. (703)305-3230

Authorized officer

Mano Padmanabhan  
Telephone No. (703)305-3900

# INTERNATIONAL SEARCH REPORT

International application No.

PCT/US03/26128

## Box I Observations where certain claims were found unsearchable (Continuation of Item 1 of first sheet)

This international report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claim Nos.:  
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☒ Claim Nos.: 14  
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:  
When read in context, it is not understood what is meant by the phrase "...writing content from the storage mechanism to". Examiner believes that some words might be missing from this claim.
3. ☐ Claim Nos.:  
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

## Box II Observations where unity of invention is lacking (Continuation of Item 2 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

☐  
☐

The additional search fees were accompanied by the applicant's protest.

No protest accompanied the payment of additional search fees.